USING PROBABILISTIC DEMAND PREDICTIONS FOR TRAFFIC FLOW MANAGEMENT DECISION SUPPORT

Anthony Masalonis[1], Sandeep Mulgund[2], Lixia Song[3], Craig Wanke[4], and Stephen Zobell[5]
The MITRE Corporation
7515 Colshire Drive, McLean, VA 22102

Abstract

This paper presents candidate information requirements and visualization concepts for explicit representation of uncertainty in decision support for air traffic flow management (TFM). Existing decision support systems for TFM provide a limited representation of uncertainty in predictions about future demand for national airspace system resources. These limitations often result in overly conservative decision-making that restricts traffic flows unnecessarily. It is believed that a better understanding of the sources and magnitude of uncertainty will assist TFM decision-makers in making timely decisions about which decisions to make and which to defer until a better understanding of a situation is developed. The research presented here is a component of a broad effort that seeks to develop risk management tools for TFM. The information presentation concepts described were developed through a set of interviews with operational experts in the TFM domain. The findings provide initial design guidance for the development of human/machine interface concepts for traffic flow management under uncertainty, and help to identify requirements for quantitative modeling of TFM uncertainty.

Introduction

Traffic flow management (TFM) is the process by which the Federal Aviation Administration (FAA), with the participation of airspace users, seeks to balance the capacity of airspace and airport resources with the demand for these resources [1]. Together with the FAA's air traffic control (ATC) function, which provides for the safe separation of aircraft from each other and from restricted areas, TFM is a central component of the nation's air traffic management (ATM) system. TFM personnel are known as Traffic Management Coordinators (TMCs) or Traffic Management Specialists (TMSs), depending on the facility in which they work. The general term for these personnel is *traffic managers*.

Unlike air traffic controllers, traffic managers do not communicate directly with pilots or ensure separation between aircraft. One of their primary responsibilities is to ensure that traffic at national airspace system (NAS) resources (e.g., airspace sectors, airports) does not exceed levels that can be safely managed by controllers. Traffic managers also endeavor to ensure fair and equitable treatment for all NAS users, i.e., operators of commercial, general aviation, military, and other aircraft. Specific problems addressed by traffic managers include congestion and weather impacting NAS resources. Specific TFM actions or *initiatives* include rerouting, metering flights to keep a certain rate of traffic flowing into a given airport or airspace, and Ground Delay Programs and Ground Stops, which delay or stop some or all traffic from or to given airport(s), for congestion or weather reasons.

Tools for TFM Decision Support

The FAA's principal decision support tool for traffic flow management is the *Enhanced Traffic Management System* (ETMS) [1]. ETMS provides real-time resource demand estimates based on predicted aircraft trajectories. In the near future, ETMS will be capable of predicting resource demand as it would be affected by proposed reroute strategies [2], and research continues towards more sophisticated strategy impact assessment capabilities [3,4].

ETMS provides several displays that help the traffic manager recognize potential traffic flow problems in advance. The *NAS Monitor* presents the alert status of all sectors in the 20 contiguous air route traffic control centers (ARTCCs) in the continental U.S., which are responsible for providing ATC services in the en route airspace between and above terminal areas. The *Center Monitor* (CM) for any ARTCC allows a traffic manager to see the predicted peak aircraft count and the alert status of all sectors within that ARTCC. Figure 1 shows a screenshot of the Center Monitor. This screenshot is from the MITRE/FAA Collaborative Routing Coordination Tools (CRCT) prototype, which was used to develop this display prior to implementation in the ETMS [4].

---

[1] Lead Systems Engineer
[2] Lead Information Systems Engineer
[3] Senior Engineer
[4] Senior Principal Engineer, Senior Member AIAA.
[5] Lead Software Systems Engineer

Each box in the CM matrix represents a 15-minute period, and the number in the box represents the maximum predicted traffic count for any single minute within that 15-minute span. The horizontal axis indicates increasing look-ahead time (LAT) (corresponding to 2000 to 2245 UTC, in this case). Each row of the matrix represents predictions for single sector (e.g. ZNY09). Next to the sector name are two sector alert thresholds (e.g. “18/18”), although currently, only one is used (which is why both thresholds are equal). This threshold is called the Monitor/Alert Parameter (MAP) and is compared to the peak count to determine whether a sector should be alerted. When the peak count is predicted to exceed the MAP for a sector, the corresponding box is colored yellow or red. Red alerts indicate that, of the aircraft involved in the peak count, enough are already airborne to exceed the MAP even if pre-departure flights are not counted. Otherwise, the alert will be yellow.

Another key tool in ETMS is the Time in Sector (TIS) display, which allows the traffic manager to analyze in detail the projected demand in any ARTCC sector. Shown in Figure 2, the TIS display is accessed by clicking on any of the cells in the CM. The TIS display shows the predicted traffic details for the selected 15-minute period for the given sector. It displays sector entry and exit times for each aircraft predicted to occupy the sector during the 15-minute period, as well as aircraft counts for each minute of the 15-minute period. As shown in the example display, it also identifies the specific aircraft that are predicted to be involved in either a yellow or red alert (or a non-alerted period)

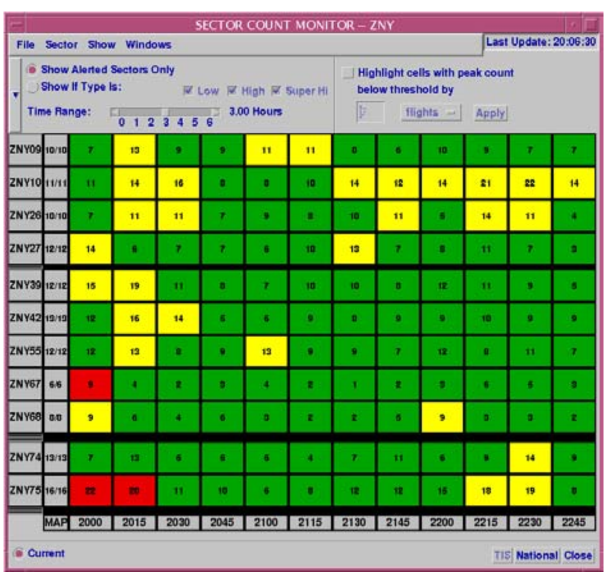

Figure 1: Center Monitor Display.

Together, these tools allow the traffic manager to monitor predicted demand for NAS resources and develop plans for dealing with potential capacity/demand imbalances. However, in their present form they provide relatively little quantification of the *uncertainty* or confidence in the information that is presented. While traffic managers understand this on some level, the lack of an explicit representation of uncertainty can manifest itself in overly conservative decision-making, which can cause unnecessary delays and sub-optimal utilization of NAS resources. This paper presents research that is part of a broader effort to manage this uncertainty through probabilistic modeling and the development of new decision support concepts for traffic flow management in the presence of uncertainty.

<u>Uncertainty in TFM Decision Support</u>

Much of the information that traffic managers require to execute their responsibilities involves estimating the demand on a resource; i.e., how many and which flights intend to use the resource. Each sector’s MAP value defines an alert threshold above which traffic management initiatives may be necessary, and airports at any given time have a defined airport arrival rate (AAR). However, the actual number of flights that can actually use a resource – i.e., its true capacity – varies with circumstances, most notably weather. At the predictive timeframes in which traffic managers operate (from 30 minutes for small localized problems out to approximately six or more hours for national-level situations), demand and capacity predictions are not certain.

Recent MITRE work [5, 6] has focused on quantifying the uncertainty existing in today’s aircraft trajectory predictions for TFM. However, today’s decision support automation for traffic management does not explicitly present the uncertainty of the demand predictions. It also does not quantify the magnitude or the likelihood of capacity reductions due to weather or other events. Although displays of weather’s location and severity are available, and capacity reductions can be inferred with some mental effort and a fair degree of subjectivity, decision support would be helpful for estimating the probabilities of both demand and capacity.

One exception to the general statement that uncertainty is not presented in today’s automation, is that sector demand predictions in the ETMS do include a crude uncertainty estimate. As discussed earlier, predicted exceedances are differentiated as yellow or red alerts depending on how many of the

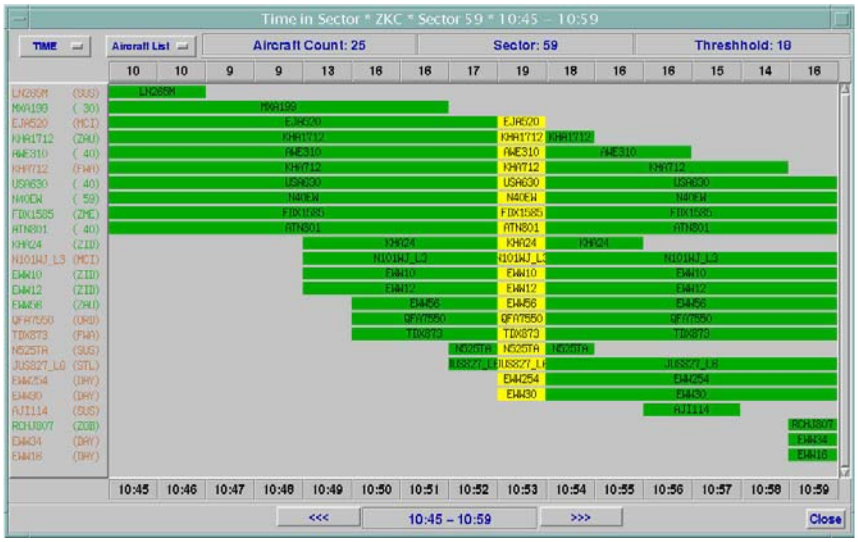

Figure 2: Time in Sector Display.

aircraft involved are airborne. This approach is based on the logic that the trajectories of aircraft not yet airborne are less certain, mainly due to the difficulty of estimating departure time. However, this rule treats departure time as the only source of uncertainty. While it is true that departure time estimation is a major source of uncertainty in demand predictions, it is not the only one. In addition, the yellow/red distinction only divides predictions into two levels of certainty, with no indication of the *degree* of that uncertainty.

Uncertainty about anticipated traffic demand makes TFM tasks more difficult in general. Therefore, one goal should be to reduce uncertainty. This can be accomplished to some extent via better data such as earlier indication from flight operators of their intended routes, by applying improved weather and traffic models, and/or by imposing order on the traffic. However, some uncertainty will always exist both in traffic demand and in capacity-reducing factors such as weather. The remaining uncertainty should be explicitly addressed in some way in strategic traffic flow management decision support, to assist problem identification and resolution.

To make effective use of demand uncertainty in TFM decision support, two major goals must be met. First, a better understanding is needed of the sources of uncertainty and the degree to which they contribute to uncertainty. Second, appropriate means must be developed for presenting information about uncertainty to the traffic manager to support real-time decision-making. Continuation of the work described in [5] and [6] is addressing the first goal. This paper describes the issues and activities involved with the second goal.

It is expected that if uncertainties can be quantitatively defined and presented to the traffic manager in an operationally useful manner, that traffic management initiatives may be more appropriate to the situation, being more likely to resolve demand/capacity imbalances without inappropriately creating inefficiency and penalizing NAS users. Decision support tools that present an explicit measure of confidence in future predictions can thus provide a foundation for *risk management* in TFM. With a clear understanding of where problem areas are likely to occur and a metric of their likelihood, a traffic manager can make informed decisions about what re-routing, delay, or other actions to perform (or not perform) and when to wait until a better prediction is available. Understanding how traffic managers can use uncertainty information in their existing decision support displays provides a step forward for the development of such risk management capabilities.

Operational Study

Purpose and Methodology

To develop insight into potential information needs for presentation of uncertainty information in TFM, a user study was conducted with former operational ATM or air carrier personnel, who are now MITRE staff. The study had four objectives:

1) Determine which information parameters regarding probabilistic demand predictions were most important for TFM decision support,

2) Derive preliminary ideas about ways to display these predictions,

3) Better understand the subjective costs of congestion and rerouting situations, and

4) Collect opinions from operational experts about the general concept of using probabilistic information for TFM decision-making.

The study investigated a number of different information parameters that could be used for traffic management decision making under uncertainty, using notional mockups to show how the information might be displayed. Structured questionnaires were used to collect *Likert* ratings [7] on the information parameters and interfaces, as well as collecting subjective cost data on various types of sector threshold exceedance situations and reroutes. General discussion sessions, using two catalyst questions, were held at the end of the study. The Likert ratings make it possible to quantify the degree to which participants agreed or disagreed with the benefits of proposed information display strategies.

Ten personnel from MITRE's Center for Advanced Aviation Systems Development (CAASD) participated. All had operational experience either as air traffic controllers or air cargo carrier dispatchers. One participant had operational experience as a traffic manager, and several were well versed in traffic flow management from their work at CAASD and elsewhere. Several participants were also former pilots.

Each session lasted approximately 90 minutes, and took place in a group setting with two to four participants attending and one or two experimenters facilitating. A brief introduction was given, participants were seated at a computer and given a packet containing the annotated user interfaces for the information parameters of interest, and traffic illustrations for the subjective cost portion of the exercise. Each participant filled out the questionnaire on his or her own computer.

Information Needs

Table 1 shows the Likert ratings on the 1-to-5 scale used for Part 1 of the questionnaire, which sought to ascertain operationally useful parameters regarding probabilistic demand predictions for TFM decision support.

The degree of MAP exceedance (i.e., count of flights in the sector above MAP) and the range of alert start times were rated highly, as was the range in peak count estimates. These results as a whole are not surprising, because sector count is the primary metric used today for TFM decision making, and the range extends its utility by explicitly acknowledging and quantifying the uncertainty existent in the particular prediction. Display of the worst case estimate by itself was not desirable. The participants' comments indicated that worst-case estimates were available from the range, so there was not a need to call out display of worst-case estimates as a unique display mode. It was also noted that displaying the worst case might result in over-restrictive decisions.

Table 1 shows that the range of possible alert start times was also deemed important. Sometimes this parameter can be uncertain due to the automation having imperfect knowledge of departure times or other factors that could affect when flights will reach the sector of interest. Representation of this element of uncertainty could help traffic managers determine the need for action.

The study results indicated a strong preference for a new color-coding scheme based on the probability of MAP exceedance, severity, and duration. This parameter was the most "popular," but also one of the most vaguely-defined ones used in the present study. It remains a challenge to define a color scheme that maps the operationally critical variables to colors. This study did not attempt to specify clear direction on how to define a color scheme, and there are many possible methods for generating a color code for a particular sector at a particular time. Due to the complexity of requirements for this information parameter, further work is definitely needed to better define a color-coding algorithm, including the collection of subjective and objective data on the

Table 1: Subject Ratings of Candidate Information Parameters

Category	Info param	Mean	Mdn	Pct. giving 4 or 5	STDEV
Peak count or occupancy	1. Best guess (today's way)	3.35	3.75	60%	0.94
	2. Range	3.70	4.00	70%	1.16
	3. Worst case	2.80	3.00	20%	0.92
	4. PDF (Probability Density Function)	3.20	3.50	50%	1.62
Alert duration	5. Best guess (today's way)	2.90	3.00	20%	0.88
	6. Range	3.30	3.00	40%	1.25
	7. Worst case	3.10	3.00	40%	0.99
Alert characteristics	8. Alert coloring system based on probability of exceeding MAP, severity & duration of exceedance	4.30	4.50	90%	0.95
	9. Range of when alert might start	3.80	4.00	80%	1.32
	10. Probability of exceeding MAP (w/o regard to by how much or for how long)	3.20	3.50	50%	1.32
Flight-specific info	11. Probabilities of various entry/exit times	2.70	3.00	30%	1.16
	12. Probability of ever entering sector	2.30	2.00	20%	1.16
	13. Route type summary (how much of predicted demand is active, proposed, early intent, etc.)	3.40	3.50	50%	1.17

utility of any proposed algorithm. Potential color coding algorithms are described briefly below, and work is underway at MITRE to define and investigate these and other algorithms.

What is clear is that the color scheme should enable at-a-glance comprehension of traffic trends within a single sector; displays supporting this comprehension have been found operationally important in past CRCT evaluations [8, 9]. The candidate displays used in the study presented peak count ranges textually (e.g., "10-13"). This may make it difficult to see trends in a single sector over time, suggesting the possible exploration of graphical sector count displays. However, were the alert coloring "smarter;" i.e., based on probability and other factors, traffic managers might be able to obtain all the trend information needed from seeing how many consecutive 15-minute time periods (the temporal resolution with which they are accustomed to working) are green, yellow, or red (or whatever sequence of colors were used).

One possible presentation is to show a *probability density function* (PDF) of the possible traffic values, indicating the likelihood of specific peak counts. Such a display is shown at middle right in Figure 10. Mixed opinions were found regarding this display. Further examination of the ratings and comments shows that this presentation evokes a "love it or hate it" response. Some operational personnel felt that the PDF represents a valuable analytical tool (e.g., the "flatness" of the curve may indicate how certain the prediction is or is not), and others feel it would cause confusion and information overload in a busy situation. Therefore, it may be that this display should be available as a drill-down function, to be accessed by the user who is interested in this level of detail and has the time to use the rich information it provides. One participant in the study indicated that intuitively, the middle value of the range would be perceived as most likely, in essence a "bell curve" shaped PDF (or some symmetrical shape with a peak in the middle, in any case). However, it is not known whether this would be the perception for all users; perhaps some would believe that all values in the range were equally likely. It is also not yet known what the shape of typical sector demand PDFs would look like in an operational environment. What is clear is that if perception did not equal reality for the particular application, there would be design implications. For example, if traffic managers did develop a miscalibrated understanding of probability from viewing the range, the PDF would become more important, because it might avert any misconception of the probabilities of various values. Although most operational personnel are not mathematicians and may not be familiar with a PDF, its intuitive nature should make it possible to train users on its operational interpretation. Therefore, training might help make this type of display more palatable to a wider range of operational users.

Uncertainties regarding specific flights were not judged to be as crucial as other "big picture" information. However, because TFM initiatives eventually require the identification of the specific flights that are contributing to the problem and that can change their route or timing to solve the problem, it would seem that this type of information is important to present to those making traffic management decisions. It may be that flight-specific information, just like the PDF, may belong in a drilldown display. ETMS's interface design follows this approach, providing all drill-down information about a particular 15-minute window within a given sector in the Time-in-Sector (TIS) display, which is considered operationally useful based on CRCT evaluations [8, 9]. A possible approach for probability-enhanced information on individual flights may be to place this information within a modified TIS interface, as discussed further below.

Another way to provide necessary support for flight-specific decisions without overloading the user is to base the specificity of the information on the degree of uncertainty. That is, the system could identify the need or likely need for initiatives at long lookahead times, when demand predictions are still uncertain, based on aggregate information, and only identify the specific flights that are to be affected by an initiative later, when uncertainty decreases and it is more certain whether and when the excess demand will occur and which flights will cause it. Analysis of traffic data is underway at MITRE to determine the degree of uncertainty in flight-specific sector entries [10]. This includes the better quantification of uncertainty in predictions about aircraft in various stages of flight (scheduled, proposed, taxi, airborne, etc.), and the investigation of the degree of uncertainty caused by a number of operational factors.

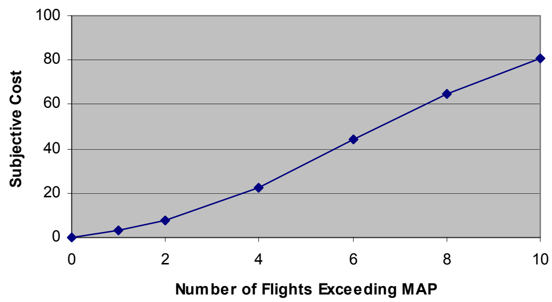

Figure 3: Subjective cost (0-100) vs. number of flights over threshold (MAP).

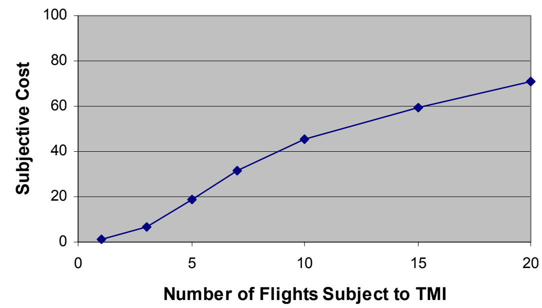

Figure 4: Subjective Cost (0-100) vs. Number of Flights affected by traffic management actions

Factors Believed to Influence Uncertainty

Useful comments were received during the interviews about specific factors that participants felt may influence demand uncertainty. These include:

- Direction and flight length (e.g., departure times, en route times, or exact routes of east-west or west-east flights across North America may be subject to greater uncertainty than short-haul or north-south flights, due to winds and total flight length),
- Degree of structure in the airspace (e.g., more structure, and thus less uncertainty, in the Eastern US than in the Midwest or West),
- Departure airport (shown by past analysis to be especially important for departure time prediction),
- Carrier (e.g., One carrier may use only one route for a given city pair, while another carrier may employ several, having more transcontinental flights and more sophisticated flight planning software)
- Knowledge of other initiatives going into effect, such as *National Playbooks*.[1]

Ongoing analysis work at MITRE will attempt to quantify the degree of uncertainty contributed by these and other factors, so that knowledge of the relevant factors for a given flight can be used to generate appropriate stochastic predictions.

Subjective Cost Ratings

Figure 3 and Figure 4 are taken from the responses to Part 2 of the questionnaire. Figure 3 shows the mean subjective cost ratings across the 10 participants, in arbitrary cost units, for each count of flights over the threshold. Figure 3 shows that on average, participants assigned almost no cost to small MAP overages, with the cost going up in a nearly linear fashion as the number of flights above the MAP increases. Figure 4 shows the mean subjective cost ratings across the 10 participants, in the same arbitrary cost units, for various numbers of flights that have to be rerouted or affected by some other traffic management initiative (TMI) such as a ground delay or an altitude change. Notice that a TMI on a few planes is not perceived as having a great cost to the NAS, but that the cost goes up slightly more sharply going from 5 to 10 flights, leveling off above 10 so that adding more planes to the TMI has less of an impact per flight.

Figure 5 and Figure 6 show the subjective costs of various threshold exceedances for various lengths of time. The two figures show two different views of the exact same data, with the axis parameters changed in order to show a different view of the data. Figure 5 indicates that the cost of a sector being at threshold, or only slightly over it, is generally believed to be quite low, even for long durations. Figure 6 shows that the cost of increasing alert duration (the steepness of each individual line) is slightly steeper when over the threshold by many planes.

Because these data come from a small population representing mostly former controllers, with less representation from the airline operations and traffic management sides, conclusions about the data are necessarily tentative. This is especially the case for the subjective cost data, which uses arbitrary units and which may vary depending on what type of operational person is making the rating (even though instructions were to consider the cost to the whole

[1] The *National Playbook* is a traffic management tool that provides the FAA and NAS users with standardized reroute strategies for handling common demand-capacity imbalance situations.

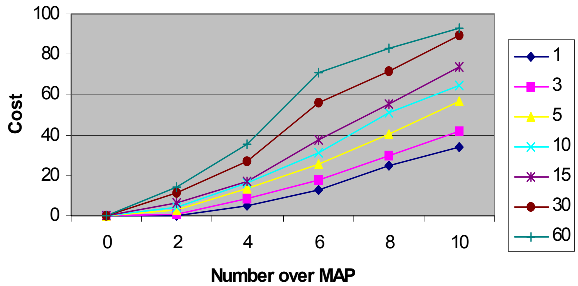

Figure 5: Subjective cost vs. MAP exceedance for different alert durations from 1 to 60 minutes.

NAS). However, subjective cost results do demonstrate that operational personnel are able to rank situations according to their severity and that decision support helping them to do so, based on subjective costs collected from a larger population and supplemented by objective data, may be possible.

It should be cautioned that due to the different priorities of all collaborators in traffic management, especially between FAA and NAS users but also among different FAA facilities and NAS user groups, it may not be possible to derive an "optimal" solution for any given situation (i.e., to reroute a given number of flights, or not to reroute). Instead, it may be preferable to present the stochastic impact of a proposed solution on multiple dimensions, such as a congestion dimension and an efficiency dimension, and allow collaborators to use this information as a starting point for decision making, applying their operational judgment to arrive at a final decision. In other words, it may not be appropriate for the automation to suggest a solution to a predicted demand imbalance that attempts to optimize on the single dimension of subjective cost.

Finally, according to the subjective cost results, alert duration by itself was considered slightly less important than magnitude of the exceedance in aircraft. Earlier studies [11, 12] showed operational utility of duration to TMCs, so it was surprising that duration did not receive higher ratings in this study. The result could be an artifact of the pictures the study participants viewed, which did not depict the duration, instead expecting participants to imagine the pictured traffic load remaining sustained for the given period of time. However, alert duration remains an operationally relevant variable: formal FAA procedures [13] dictate prioritization of alerts predicted to last more than five minutes. This might be handled from a user interface perspective by giving such alerts greater prominence when the predicted duration is longer than five minutes. The somewhat surprising findings regarding alert duration

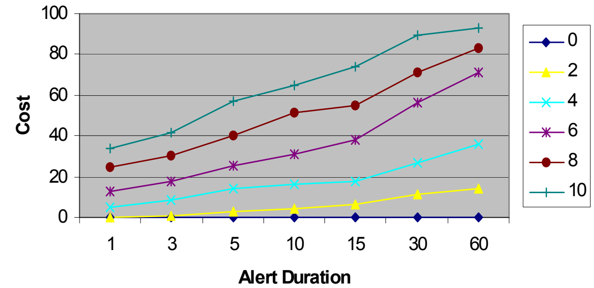

Figure 6: Subjective cost vs. alert duration for different MAP exceedance values from 0 to 10.

suggest a more rigorous investigation of the significance of this information parameter.

<u>Application to TFM Decision Support</u>

The results of the interviews with subject matter experts described above provide a foundation for the development of candidate decision support concepts for probabilistic TFM decision support. Here we synthesize these results and present preliminary display concepts. The current research focuses on the modeling and visualization of probabilistic measures of airspace demand. Therefore, attention focuses on the development of display concepts for achieving this understanding. The development of tools for "what if" analyses of rerouting and other traffic management strategies is a focus of ongoing and future work.

<u>Color Coding Strategies</u>

The results of the operational study indicate that users would like to see a color-coding scheme based on probability of MAP exceedance, alert severity (i.e., number of aircraft over MAP), alert and duration. Such a scheme would replace the color coding used on existing TFM displays, which is based on deterministic predictions of sector occupancy and the status (airborne or not) of aircraft causing an alert condition. Several of the study participants indicated that a three-color scheme would be desirable. This lends itself to the canonical green/yellow/red demarcation, but the challenge is how to map those colors to relevant underlying variables. Any candidate color coding scheme should be easy to understand, objective, and tied to operationally relevant variables.

Figure 7 presents one possible color coding methodology. In this design, color coding is based strictly on the probability of MAP exceedance in the given sector/time window. The probability thresholds shown are for illustration purposes only; in practical implementation it would be necessary to make

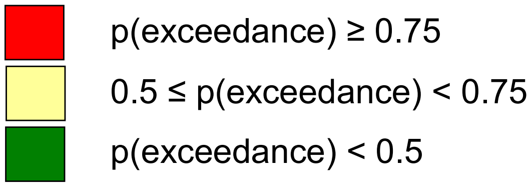

Figure 7: Center Monitor Color Coding based on MAP Exceedance Probability

educated decisions about where the boundaries between green, yellow, and red should be paced for maximum operational effectiveness. For the purposes of discussion, a cell in the CM would be shown in green if there is less than a 0.5 probability that an alert will occur. Color coding changes to yellow for probabilities between 0.5 and 0.75, and is red beyond a 0.75 probability. The benefit of this color coding scheme is that it is relatively simple and objective: the color of a particular cell ties directly to a single computed quantity that is straightforward to understand. However, a drawback is that the raw probability measures may require additional mental synthesis to be transformed into quantities used directly by traffic flow managers for decision-making. For example, the probability-based color-coding shown does not explicitly reflect measures of alert severity or duration, although the probabilities themselves are statistically-related to those factors.

Figure 8 presents another color coding scheme that attempts to remedy this limitation. In this case, sector color coding is based on the *expected cost* to the NAS for the predicted alert condition. The expected cost *EC* is defined as:

$$EC = p(alert) \cdot Cost(alert) \qquad (1)$$

where $p(alert)$ is the probability of a given alert condition and $Cost(alert)$ is the "cost" to the NAS of that alert condition (that is, the cost of having the given number of aircraft over the MAP for the given amount of time). The cost would be based on cost curves identified during the interviews with operational personnel, eventually augmented with further interviews using a larger sample, and/or other techniques such as modeling. Hypothetical cost curves based on the present data are shown here.

As an example, if the alert duration were predicted to be 60 minutes, with average MAP exceedance (or however MAP exceedance were reduced to a single number) being 6, this would map to a

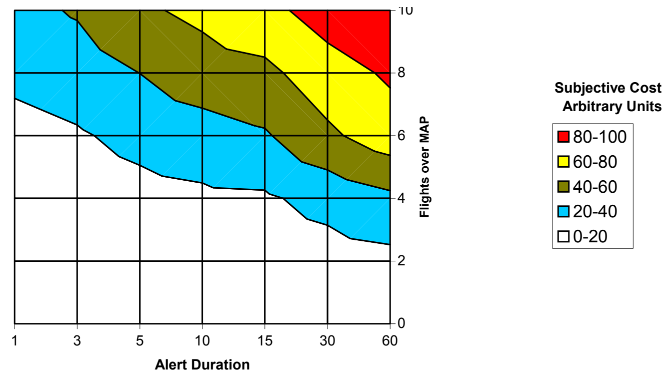

Figure 8: Conceptual Color Coding based on Expected Cost to National Airspace System.

cost corresponding to a yellow alert condition. In Figure 8, five colors are used in order to illustrate that the ideal number of colors is still an open question. The number of regions could be limited to three, as in the color coding technique used earlier in Figure 7. Regardless of the number of colors, the benefit of the color coding scheme in Figure 8 is that it may be tied to variables that are more operationally relevant than the simple probability of alert, such as its severity and duration. Its drawback, as mentioned earlier, is that the cost curves are subjective and may vary between NAS users and traffic management facilities, as well as between different NAS user facilities, between different traffic management facilities, and even between different personnel within a facility.

These examples provide two possible approaches to color coding of probabilistic TFM displays. Further research and investigation is needed to establish which schemes offer the most potential for prototyping and simulation-based evaluation.

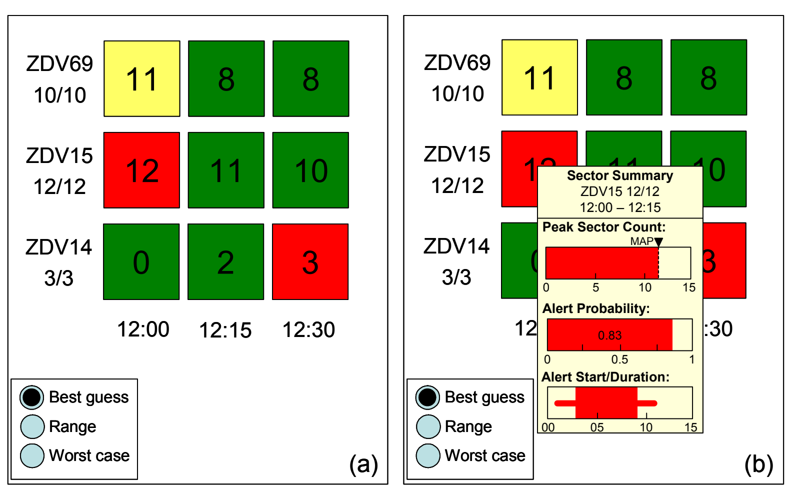

Figure 9: Hypothetical Probabilistic Center Monitor: (a) Basic Display. (b) Sector Summary via Mouse Rollover.

Probabilistic Center Monitor

Once a color coding scheme is selected, we can apply it to develop a version of the CM that presents probabilistic predictions of sector occupancy within an ARTCC. Figure 9 shows a hypothetical example of such a display. On the left is the basic display; each cell is colored using whatever probability-based scheme is eventually implemented. The user will have the option to display the "best guess" of the sector count, or switch to a display mode that presents the 95% confidence interval (or some appropriate interval) in the range of sector count predictions or the worst case prediction. As discussed earlier, the Worst Case mode may not be necessary, since the Range mode will provide that information as well.

When the user positions the mouse cursor over a particular cell (without clicking on the cell), a "rollover" display will be presented to provide preliminary drill-down information on the alert predicted in that sector/time. Figure 9b presents an example of how this rollover display might be formatted. The Sector Summary rollover shows bar graph representations of the predicted peak sector count, the alert probability, and a box-and-whisker representation of the range in the predicted beginning and end of the alert. The leading line presents a graphical representation of the range in the predicted start time of the alert, while the tail end shows the range in the predicted end time. All bar graphs are color coded using the same logic as that driving the CM. If an alert is not predicted (i.e., green), the bottom-most graph would not be shown.

Probabilistic Time-in-Sector Display

The final component of the envisioned probabilistic TFM displays is an enhanced version of the existing CRCT/ETMS TIS display, which provides detailed information about a particular sector in a given 15-minute time window. The intention of this enhanced display is to collect within one window *all* of the information that a traffic manager may wish to see about a particular sector/time from the CM. Figure 10 presents a possible design for such a display.

The content is largely the same as on the existing TIS display, however some candidate additions have been made to provide better information integration as well as enhancements for visualization of probabilistic information. All of this information would be presented within a single window, to co-locate all of the information relevant to this sector and 15-minute time window. On the top right corner is a duplicate of the Sector Summary display, which would provide the same content as the rollover

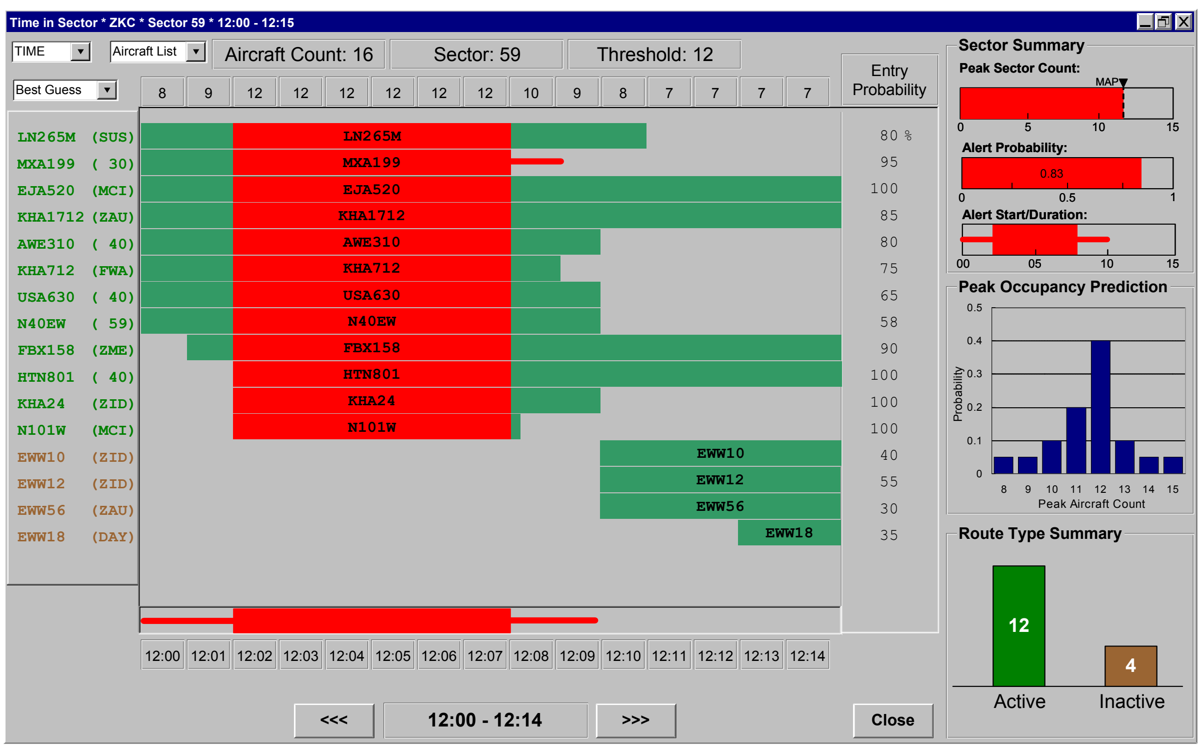

Figure 10. Hypothetical Probabilistic Time-in-Sector Display.

shown earlier in Figure 9b. Below it might be a rendering of the probability density function of sector occupancy for the time window of interest. At the bottom would be a route type summary display for the sector, which allows the traffic manager to ascertain at a glance which of the predicted flights are already airborne, which have filed plan information, and so on – another information parameter deemed important by study participants and known to affect uncertainty [6]. The stackbar of inactive flights separates them based on current status: filed, scheduled, etc.

The bulk of the window would be taken up the standard TIS display, with color coding made consistent with the remainder of the PTDV system. At the bottom would appear the predicted start/end times of the alert condition. This information would be presented via the same kind of box-and-whisker representation shown in the sector summary rollover, which allows the user to see the range in the predicted start and end of the alert.

TFM Uncertainty and Risk Management

The work described in this paper has helped to provide insight into potential information needs and data visualization concepts for presentation of quantitative uncertainty information in traffic flow management. These concepts focus on helping the traffic manager develop an understanding of the likely evolution of *demand* for NAS resources within a time period of interest. Predictions of the *capacity* of those resources hours into the future will also be uncertain, especially on severe weather days. MITRE and other organizations are researching probabilistic forecast models of NAS capacity.

Two key factors that influence capacity are weather and traffic complexity. Effects of weather will be determined using new forecasting techniques that are tailored for assessing the impact on aviation. These will necessarily be probabilistic, due to the difficulty of predicting weather with perfect accuracy. Complex traffic flow patterns involving heavy crossing or nonstandard routing (which often is a result of weather rerouting) are associated with higher perceived workload, and thus can possibly reduce the amount of traffic that can be handled [11, 12]. Predicting complexity, like predicting traffic count and weather, is possible at longer lookahead times but is more uncertain [12].

Probabilistic capacity forecasts will be combined with probabilistic demand forecasts to form probabilistic forecasts of NAS congestion, i.e., demand-capacity imbalances. These will present the risks of future weather and congestion problems. Using such information, traffic managers must formulate and implement plans for addressing potential congestion or weather-induced reroute situations. For any given level of uncertainty, a traffic manager will have to choose between taking action now and waiting to develop a better understanding of the situation before deciding what to do. There will be risks associated with either course of action: reroutes, altitude changes, or delays that are implemented in the presence of too much uncertainty may prove to be inadequate or excessive, while waiting for too long before enacting some strategy may reduce the traffic manager's flexibility in dealing with a problem situation. These considerations motivate the exploration of effective *risk management* strategies for TFM, in which uncertainty representations about predicted demand for NAS resources are coupled with decision automation technologies that help the traffic manager understand the potential risks of action or inaction, and assist in taking appropriate actions in such a way that balance these competing considerations. Ongoing work at MITRE is exploring how to leverage quantitative models of NAS resource demand uncertainty [5, 10] with effective visualization and decision support designs to develop risk management models for TFM.

Conclusions

This paper has presented an operational study that was conducted to develop insight into the potential information needs of air traffic flow managers for presentation of uncertainty information. The findings provide initial design guidance for the development of human/machine interface concepts for traffic flow management under uncertainty. These designs must be refined based on feedback from a broader user community and made congruent with anticipated operational concepts for TFM risk management. Simulation-based human-in-the-loop experimentation will make it possible to evaluate the benefits of displaying uncertainty information on a traffic manager's *situation awareness* [14] and decision-making effectiveness. The work described here concerned identifying the information requirements for characterizing the uncertainty in predicted demand for NAS resources. Future work will entail integrating such information with capacity uncertainty. Future work will also involve studying the application of advanced risk management and automation design concepts that not only present the future predicted state, but also assist the traffic manager and other stakeholders in choosing an appropriate course of action.

<u>References</u>

1. *Enhanced Traffic Management System (ETMS) Reference Manual*, Version 7.6, Report No. VNTSC-DTS56-TMS-004, Volpe National Transportation Systems Center, U.S. Department of Transportation, May 2003.

2. Wanke, C, L. Carlson, J. DeArmon, S. Mayer, S. Miller, and G. Schultz, "New Tools for Strategic Route Planning: Collaborative Routing Coordination Tools," *Handbook of Airline Strategy*, edited by G. Butler and M. Keller, The McGraw-Hill Companies, New York, 2001.

3. Wanke, C., D. Berry, J. DeArmon, S. Miller, N. Taber, L. Rhodes, and M. Yee, "Decision Support for Complex Traffic Flow Management Actions: Integrated Impact Assessment," AIAA-2002-5015, *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Monterey, CA, 5-8 August 2002.

4. Sridhar, B., G. Chatterji, S. Grabbe, and K. Sheth, "Integration of Traffic Flow Management Decisions," AIAA-2002-3748, *AIAA Guidance, Navigation, and Control Conference and Exhibit*, Monterey, CA, 5-8 August 2002.

5. Wanke, C., Callaham, M., Greenbaum, D., & Masalonis, A., "Measuring Uncertainty In Airspace Demand Predictions For Traffic Flow Management Applications," *Proceedings of the 2003 Guidance, Navigation, and Control Conference*, Austin, TX (August).

6. Lindsay, K.S., Greenbaum, D.P., & Wanke, C.R. (2003). "Impact of pre-departure uncertainty on Collaborative Routing Coordination Tools (CRCT) prediction performance. " *Presented at American Institute of Aeronautics and Astronautics (AIAA) Aircraft Technology, Integration, and Operations (ATIO) Conference*.

7. Rubin, J. (1994). *Handbook of Usability Testing*, New York: John Wiley & Sons.

8. Carlson, L.S., Kapoor, P.K., & Rhodes, L.R. (1999). *Field Evaluation Report for ZKC Stand Alone CRCT: Phase 1 Results*. Rept. No. MP99W053. McLean, VA: The MITRE Corporation.

9. Masalonis, A.J., Connolly, K.A., Hunt, R.A., Edmondson, S.M., Jacobs, G.J., Berry, D.L., Simons, E.M., Kapoor, P.K., & Hawkins, P.S. (2001). *Operational Evaluation Report (ZKC, ZID, ATCSCC)*. Rept. No. MP01W181. McLean, VA: The MITRE Corporation.

10. Wanke, C., Mulgund, S., Greenbaum, D., Song, L. (2003). "Modeling Traffic Prediction Uncertainty for Traffic Management Decision Support," *submitted to the 2004 Guidance, Navigation, and Control Conference.*

11. Masalonis, A.J., Callaham, M.B., & Wanke, C.R. (2003). "Dynamic Density and complexity metrics for realtime Traffic Flow Management." *Presented at 5th USA/Europe Air Traffic Management R&D Seminar*, Budapest, Hungary.

12. Masalonis, A.J., Callaham, M.B., Figueroa, Y., & Wanke, C.R. (2003). "Indicators of airspace complexity for Traffic Flow Management decision support." *Presented at 12th International Symposium on Aviation Psychology*, Dayton, OH.

13. Federal Aviation Administration. *Air Traffic Control*. Order No. 7110.65. Washington, DC.

14. Endsley, M.R. (1995). "Toward a theory of Situation Awareness in dynamic systems." *Human Factors*, 35(1), 32-45.